

# DiVinE-CUDA – A Tool for GPU Accelerated LTL Model Checking\*

Jiří Barnat

Faculty of Informatics  
Masaryk University  
Brno, Czech Republic  
barnat@fi.muni.cz

Luboš Brim

Faculty of Informatics  
Masaryk University  
Brno, Czech Republic  
brim@fi.muni.cz

Milan Česka

Faculty of Informatics  
Masaryk University  
Brno, Czech Republic  
xceska@fi.muni.cz

In this paper we present a tool that performs CUDA accelerated LTL Model Checking. The tool exploits parallel algorithm MAP adjusted to the NVIDIA CUDA architecture in order to efficiently detect the presence of accepting cycles in a directed graph. Accepting cycle detection is the core algorithmic procedure in automata-based LTL Model Checking. We demonstrate that the tool outperforms non-accelerated version of the algorithm and we discuss where the limits of the tool are and what we intend to do in the future to avoid them.

## 1 Introduction

Verification and validation became an important part of the design process. Unfortunately, the gap between the complexity of systems the current formal verification tools can handle and the complexity of systems built in practice is still quite wide. Therefore, any technique that accelerates the verification process is highly desirable. A possible way to reduce the delay due to the formal verification process is to accelerate the computation of verification tools using contemporary parallel hardware. Hardware platforms such as multi-core multi-cpu systems or many-core hardware accelerators, e.g. GPGPUs, have recently received a lot of attention in this aspect.

CUDA (Compute Unified Device Architecture) is a parallel computing architecture developed by NVIDIA [7]. Recently, it has been successfully used to accelerate formal verification process for selected settings. In [4] authors demonstrated significant speedup in the verification of probabilistic systems, while in [8, 9] CUDA has been used to accelerate disk-based model checking and state space generation. Let alone the CUDA technology, other many-core hardware acceleration platforms have been tried. For example, an implementation of FPGA accelerated Mur $\phi$ [13] verification tool has been reported in [10].

In this paper we introduce a new CUDA accelerated verification tool for model checking formulas of Linear Temporal Logic (LTL). The problem of LTL model checking is well established problem in the formal verification community. Computationally the problem reduces to the problem of detection of an accepting cycle in a directed graph [14]. The new tool builds upon the DiVinE [2] framework, hence the name of the tool is *DiVinE CUDA*.

## 2 DiVinE CUDA Algorithmics

DiVinE-CUDA employs algorithm MAP [5] for accepting cycle detection. The algorithm is, however, formulated as a repeated matrix-vector product procedure [3] in order to efficiently utilize CUDA archi-

---

\*This work has been partially supported by Czech Science Foundation grants No. 201/09/P497, 102/09/H042 and by Academy of Sciences of Czech Republic grant No. 1ET408050503.

ture. The idea of the MAP algorithm is as follows. Given a directed graph with accepting vertices, the algorithm impose ordering on accepting vertices and repeatedly computes the maximal (w.r.t. the ordering) accepting predecessor  $map(u)$  for every accepting vertex  $u$  in the graph. If the algorithm detects an accepting vertex that is its own maximal accepting predecessor, then the vertex lies on an accepting cycle and the algorithm terminates. In the other case, all accepting vertices that were maximal accepting predecessors for some other vertices are marked as non-accepting (because they do not lie on an accepting cycle) and the procedure is restarted (goes to the next iteration). The algorithm terminates either if accepting cycle is found or there are no more accepting vertices in the graph. From technical reasons we employ MAP algorithm on a transposed state space graph, note that graph transposition preserves the presence of accepting cycles.

The main computation demanding step of the algorithm is the computation of the maximal accepting predecessor for every accepting vertex. This is done by means of value propagation of accepting vertices along edges in the graph. If multiple values are propagated into a single vertex, the maximum among all the incoming values and the value of the vertex is computed and used for further propagation. Every vertex keeps the maximum value that has been propagated through the vertex. Once a fix-point is reached (no value can be improved), values of maximal accepting successors are computed.

In DiVinE CUDA tool it is the maximal accepting successor computation that is accelerated with CUDA device. In particular, relevant parts of the graph to be analyzed are represented in an adjacency matrix. Having the matrix, the value propagation can be realized as matrix-vector product [3] for computation of which the CUDA architecture is known to be extremely efficient [11].

When initiated the DiVinE CUDA tool proceeds as follows. It starts a thread that computes the adjacency matrix needed for CUDA processing. We use CSR (compressed sparse row) format to store the matrix. Note that we do not list all reachable states in the matrix, but only those that are in components containing some accepting vertices [12]. This feature significantly reduces the size of the matrix to be handled. (The size reduction is up to 20-30% of the full size in most cases). At the same time the tool runs a second thread that repeatedly performs CUDA accelerated accepting cycle detection on the part of the matrix that has been computed so far. If an accepting cycle is present in that part of the graph it is discovered before the full state space is generated. Therefore, DiVinE CUDA works on-the-fly.

### 3 Using the Tool

DiVinE-CUDA is a tool that stems from parallel and distributed LTL Model Checker DiVinE [2, 1]. As such, DiVinE-CUDA tool uses DiVinE native modeling language DVE [2]. In DVE modeling language the system to be verified is given as an asynchronous network of communicating finite automata. Transitions of every automaton in the network can be augmented with guards, buffered and unbuffered channel communication primitives, and variable updates.

The scheme of how the DiVinE CUDA tool should be used is given in Figure 1. Having prepared the model either directly as a *.dve* file or from a *.mdve* template using `divine.preprocessor` the user has to specify the property to be verified. The property can be given either directly as a property automaton (also known as never claim automaton) in the model file, or as (a set of) LTL formula(s) in a separate file, in which case the files have to be further processed by `divine.combine` tool to get a model file with the property automaton.

The next step in the verification process is to produce precompiled version of the model using `divine.precompile` tool. Precompiled version of the model (file with extension *.dveC*) is actually a dynamically linked library containing functions to generate states of the model with specification. Fi-

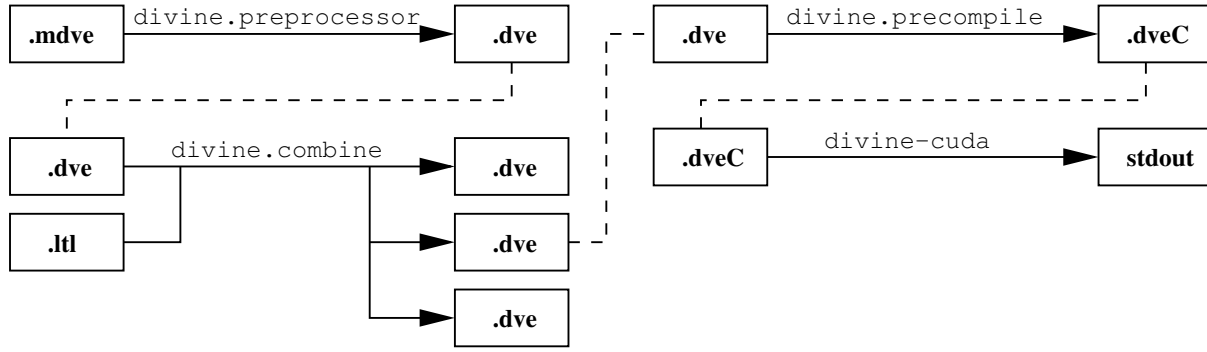


Figure 1: DiVinE CUDA work-flow.

nally, the precompiled representation of the model is used as an input for the `divine-cuda` tool itself.

During the computation the tool reports periodically the numbers of generated states and transitions, numbers of MAP iterations and CUDA device calls made so far to the standard output. At the end the tool outputs whether an accepting cycle has been found, in which case the given model does not satisfy the specification, or whether no accepting cycle has been discovered, i.e. the specification is satisfied.

## 4 Experiments

To briefly evaluate our tool we compared our implementation of CUDA accelerated MAP algorithm with the existing algorithms implemented in the DiVinE-Cluster version 0.8.2 model checker. For the comparison we used selected DiVinE native models including leader election protocol, elevator cabin system, Peterson’s and Anderson’s solutions to mutual exclusion problem and dining philosophers. We tested both the models with specification error (with an accepting cycle) as well as models without a specification error. All the experiments were run on a Linux workstation equipped with two AMD Phenom(tm) II X4 940 Processors @ 3MHz, 8 GB DDR2 @ 1066 MHz RAM and NVIDIA GeForce GTX 280 GPU with 1GB of GPU memory.

Table 1 provides details on run-times of individual algorithm parts. As for the CUDA MAP algorithm, the total run-time includes the initialization time (not reported in the table), CSR construction time (*CSR time*), and time spent on CUDA computation (*CUDA time*). Note that the first iteration of CPU MAP is actually slower than construction of the CSR representation. This is because the first iteration of the CPU MAP not only generates the state space, but also computes first stable values of *map*. Just for curiosity we also compare the performance of the new tool with DiVinE Cluster tool running OWCTY Algorithm [6]. Algorithms MAP and OWCTY were running on a single core.

Table 2 gives a comparison of overall run-times for both valid and invalid model checking instances. Though, the overall speedup is not that significant, it is still impressive. We can also see that the burden of data preparation is huge compared to the CUDA processing itself.

## 5 Availability and Future Work

At the moment the tool cannot handle models for which the corresponding *reduced* matrix of the graph does not fit the memory of a single CUDA device, it lacks the ability of counterexample generation, and cannot employ multiple threads to compute the CSR representation in parallel. We intend to address all

Model	<i>accepting cycle</i>	CUDA MAP			CPU MAP				CPU OWCTY	
		<i>CSR time</i>	<i>CUDA time</i>	<i>total time</i>	<i>1st iter. time</i>	<i>other iter. time</i>	<i>total time</i>	<i># iter.</i>	<i>reachability time</i>	<i>total time</i>
elevator 1	N	26	7	34	44	56	100	16	24	41
leader	N	87	1	90	97	600	697	17	90	297
peterson 1	N	105	6	113	175	270	445	16	110	188
anderson	N	31	7	39	64	51	115	5	33	113
elevator 2	Y	33	1	35	50	–	50	1	41	177
phils	Y	45	1	47	295	102	397	5	180	576
peterson 2	Y	25	5	31	173	–	173	1	114	404
bakery	Y	24	1	26	240	–	240	1	219	907

Table 1: Comparison of run-times (in seconds) for CUDA accelerated MAP algorithm, non-accelerated MAP algorithm and OWCTY algorithm.

Models	CUDA MAP	CPU MAP		CPU OWCTY	
	<i>total time</i>	<i>total time</i>	<i>CUDA MAP speedup</i>	<i>total time</i>	<i>CUDA MAP speedup</i>
non-accepting	276	1357	4.92	639	2.32
accepting	139	860	6.19	2064	14.87
both	415	2173	5.24	2730	6.51

Table 2: The overall run-times in seconds, and speedup of the whole model checking procedure.

	CUDA MAP			CPU MAP	CPU OWCTY
	<i>CSR time</i>	<i>CUDA time</i>	<i>total time</i>		
<b>1 core:</b>	386	+ 29 =	415	Total time: 2 173 Speedup: <b>5.24</b>	Total time: 2 730 Speedup: <b>6.51</b>
<b>2 cores:</b>	193	+ 29 =	222	Total time: 1087 Speedup: <b>4.87</b>	Total time: 1365 Speedup: <b>6.15</b>
<b>4 cores:</b>	97	+ 29 =	126	Total time: 544 Speedup: <b>4.32</b>	Total time: 683 Speedup: <b>5.42</b>
<b>8 cores:</b>	49	+ 29 =	78	Total time: 272 Speedup: <b>3.48</b>	Total time: 342 Speedup: <b>4.38</b>

Table 3: A hypothetical speedup of DiVinE CUDA w.r.t. multicore parallel algorithms. We suppose optimal (linear) speed-up for both parallel algorithms MAP and OWCTY and for the CSR construction phase of the CUDA MAP algorithm.

these issues in the next version of the tool. As for the run-times, we expect significant improvement due to parallel preparation of the CSR graph representation. See Table 3. As for the limit on the size of the verification problem, we plan to introduce sort of clever swapping mechanism of the matrix stored in the GPU memory and to extend the memory available by employ multiple CUDA devices.

DiVinE CUDA tool is freely available from DiVinE web pages <sup>1</sup> where we provide both download and install instructions as well as simple tutorial on using the tool.

## References

- [1] J. Barnat, L. Brim & P. Ročkal (2008): *DiVinE Multi-Core – A Parallel LTL Model-Checker*. In: *Automated Technology for Verification and Analysis*, LNCS 5311. Springer, pp. 234–239.
- [2] J. Barnat, L. Brim, I. Černá, P. Moravec, P. Ročkal & P. Šimeček (2006): *DiVinE – A Tool for Distributed Verification (Tool Paper)*. In: *Computer Aided Verification (CAV 2006)*, LNCS 4144. Springer, pp. 278–281.
- [3] J. Barnat, L. Brim, M. Češka & T. Lamr (2009): *CUDA accelerated LTL Model Checking*. Technical Report FIMU-RS-2009-05, Faculty of Informatics, Masaryk University.
- [4] D. Bosnacki, S. Edelkamp & D. Sulewski (2009): *Efficient Probabilistic Model Checking on General Purpose Graphics Processors*. In: *Model Checking Software (SPIN 2009)*, LNCS 5578. Springer, pp. 32–49.
- [5] L. Brim, I. Černá, P. Moravec & J. Šimša (2004): *Accepting Predecessors are Better than Back Edges in Distributed LTL Model-Checking*. In: *Formal Methods in Computer-Aided Design (FMCAD’04)*, LNCS 3312. Springer, pp. 352–366.
- [6] I. Černá & R. Pelánek (2003): *Distributed Explicit Fair Cycle Detection (Set Based Approach)*. In: *Model Checking Software (SPIN’03)*, LNCS 2648. Springer, pp. 49–73.
- [7] NVIDIA CUDA Compute Unified Device Architecture. [http://www.nvidia.com/object/cuda\\_home.html](http://www.nvidia.com/object/cuda_home.html), August 2009.
- [8] S. Edelkamp & D. Sulewski: *Technical Report on GPU based Model Checking*. [http://www.tzi.de/~edelkamp/GPU\\_Technical.pdf](http://www.tzi.de/~edelkamp/GPU_Technical.pdf), August 2008.
- [9] Stefan Edelkamp & Damian Sulewski (2009): *Parallel State Space Search on the GPU*. In: *International Symposium on Combinatorial Search (SoCS 2009)*.
- [10] M. E. Fuess, M. Leeser & T. Leonard (2008): *An FPGA Implementation of Explicit-State Model Checking*. In: *16th IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM 2008)*. IEEE Computer Society, pp. 119–126.
- [11] Michael Garland (2008): *Sparse Matrix Computations on Manycore GPU’s*. In: *Proceedings of the 45th annual conference on Design automation (DAC’08)*. ACM, pp. 2–6.
- [12] A. L. Lafuente (2002): *Simplified distributed LTL model checking by localizing cycles*. Technical Report 00176, Institut für Informatik, University Freiburg, Germany.
- [13] U. Stern & D. L. Dill (1997): *Parallelizing the Murφ Verifier*. In: O. Grumberg, editor: *Proceedings of Computer Aided Verification (CAV ’97)*, LNCS 1254. Springer-Verlag, pp. 256–267.
- [14] M. Y. Vardi & P. Wolper (1986): *An Automata-Theoretic Approach to Automatic Program Verification (Preliminary Report)*. In: *Proceedings 1st Annual IEEE Symp. on Logic in Computer Science, LICS’86*, Cambridge, MA, USA, 16–18 June 1986. IEEE Computer Society Press, Washington, DC, pp. 332–344.

---

<sup>1</sup><http://divine.fi.muni.cz/page.php?page=divine-cuda>